

**Kaufman -
Trading
Systems &
Methods -**

Chap02 4-13

VISIT...

LANZAROTE
Caliente.com

Although this is not exact because of the use of average prices for intervals, it does closely represent the average price relative to time. There are two other averages for which time is an important element: the geometric and harmonic means.

Geometric Mean

The geometric mean represents a growth function in which a price change from 50 to 100 is as important as a change from 100 to 200.

$$G = (a_1 \times a_2 \times a_3 \times \cdots \times a_n)^{(1/n)} \text{ or } (@\text{mult}(\text{list}))^{(1/\text{length of list})}$$

To solve this mathematically, rather than using a spreadsheet, the preceding equation can be changed to either of two forms:

$$\ln G = \frac{\ln a_1 + \ln a_2 + \cdots + \ln a_n}{n}$$

or

$$\ln G = \frac{\ln (a_1 \times a_2 \times a_3 \times \cdots \times a_n)}{n}$$

The two solutions are equivalent. Using the price levels in Table 21, disregarding the time intervals, and substituting into the first equation:

$$\ln G = \frac{\ln 40 + \ln 80 + \ln 120 + \ln 160 + \ln 200}{5}$$

Had one of the periods been a loss, that value would simply be negative. We now perform the arithmetic to solve the equation:

$$\ln G = \frac{3.689 + 4.382 + 4.787 + 5.075 + 5.298}{5}$$

$$\ln G = 4.6462$$

$$G = 104.19$$

The geometric mean has advantages in application to economics and prices. A classic example is to compare a tenfold rise in price from 100 to 1,000 with a fall to one-tenth from 100 to 10. An arithmetic mean of 10 and 1,000 is 505, while the geometric mean gives

$$G = (10 \times 1000)^{(1/2)} = 100$$

which shows the relative distribution as a function of comparable growth. Due to this property the geometric mean is the best choice when averaging ratios that can be either fractions or percentages.

Quadratic Mean

The quadratic mean is as calculated:

$$Q = \sqrt{\frac{\sum a^2}{N}} \text{ or } (@\text{sqrt}((@\text{sum}(\text{list of } a^2) \text{ length of list}))$$

The square root of the mean of the square of the items (rootmeansquare) is most well known as the basis for the standard deviation. This will be discussed later. in the section "Dispersion and Skewness."

Harmonic Mean

The harmonic mean is more of a timeweighted average, not biased toward higher or lower values as in the geometric mean. A simple example is to consider the average speed of a car that travels 4 miles at 20 mph, then 4 miles at 30 mph. An arithmetic mean would result in 25 mph, without

considering that 12 minutes were spent at 20 mph and 8 minutes at 30 mph. The weighted average would give

$$W = \frac{(12 \times 20) + (8 \times 30)}{12 + 8} = 24$$

The harmonic mean is

$$\frac{1}{H} = \frac{\frac{1}{a_1} + \frac{1}{a_2} + \dots + \frac{1}{a_n}}{n}$$

which can also be expressed as

$$H = n / \sum_{i=1}^n \left(\frac{1}{a_i} \right) \text{ or length of list} / \text{sum(list of } 1/a)$$

For two or three elements, the simpler forms can be used:

$$H_2 = \frac{2ab}{a+b} \quad H_3 = \frac{3abc}{ab+ac+bc}$$

This allows the solution pattern to be seen. For the 20 and 30 mph rates of speed, the solution is

$$H_2 = \frac{2 \times 20 \times 30}{20 + 30} = 24$$

which is the same answer as the weighted average. Considering the original set of numbers again, the basic form of harmonic mean can be applied:

$$\begin{aligned} \frac{1}{H} &= \frac{\frac{1}{40} + \frac{1}{80} + \frac{1}{120} + \frac{1}{160} + \frac{1}{200}}{5} \\ &= \frac{.05708}{5} = .01142 \end{aligned}$$

$$H = 87.59$$

We might apply the harmonic mean to price swings, in which the first swing moved 20 points over 12 days, and the second swing moved 30 points over 8 days.

DISTRIBUTION

The measurement of distribution is very important because it tells you generally what to expect. We cannot know what tomorrow's S&P trading range will be, but we have a high level of confidence that it will fall between 300 and 800 points. We have a slightly lower confidence that it will vary from 400 to 600 points. We have virtually no chance of picking the exact range. The following measurements of distribution allow you to put a value on the chance of an event occurring.

14

Frequency Distributions

The frequency distribution can give a good picture of the characteristics of the data. To know how often sugar prices were at different price levels, divide prices into 10 increments (e.g., 5.01 to 6.00, 6.01 to 7.00, etc.), and count the number of times that prices fall into each interval. The result will be a distribution of prices as shown in Figure 22. It should be expected that the distribution of prices for a physical commodity interest rates (yield). or index markets, will be skewed toward the lefthand side (lower prices or yields) and have a long tail toward higher prices on the righthand side. This is because prices remain at higher levels for only a

short time relative to their longterm characteristics. Commodity prices tend to be bounded on the lower end, limited in their downside movement. by production costs and resistance of the suppliers to sell at prices that represent a loss. On the higher end, there is not such a clear point of limitation; therefore, prices move much further up during periods of extreme shortage relative to demand.

The measures of central tendency discussed in the previous section are used to qualify the shape and extremes of price movement shown in the frequency distribution. The general relationship between the results when using the three principal means is


arithmetic mean > geometric mean > harmonic mean

Median and Mode

Two other measurements, the median and the mode, are often used to define distribution. The median, or middle item, is helpful for establishing the center of the data: it halves the number of data items. The median has the advantage of discounting extreme values. which might distort the arithmetic mean. The mode is the most commonly occurring value in Figure 23 the mode is the highest point.


In a normally distributed price series, the mean, median, and mode will all occur at the same value; however, as the data become skewed, these values will move farther apart. The general relationship is:

FIGURE 2-2 Hypothetical price-frequency distribution.



15

FIGURE 23 Hypothetical price distribution skewed to the right, showing the relationship of the mode, median, and mean.



mean > median > mode

The mean, median, and mode help to tell whether data is normally distributed or skewed. A normal distribution is commonly called a bell curve, and values fall equally on both sides of the mean. For much of the work done with price and performance data, the distributions tend to extend out toward the right (positive values) and be more cut off on the left (negative values). If you were to chart a distribution of trading profits and losses based on a trend system with a fixed stoploss, you would get profits that could range from zero to very large values, while the losses would be theoretically limited to the size of the stop. Skewed distributions will be important when we try to measure the probabilities later in this chapter.

Characteristics of the Principal Averages

Each averaging method has its unique meaning and usefulness. The following summary points out their principal characteristics:

The arithmetic mean is affected by each data element equally, but it has a tendency to emphasize extreme values more than other methods. It is easily calculated and is subject to algebraic manipulation.

The geometric mean gives less weight to extreme variations than the arithmetic mean and is most important when using data representing ratios or rates of change. It cannot always be used for a combination of positive and negative numbers and is also subject to algebraic manipulation.

The harmonic mean is most applicable to time changes and, along with the geometric mean, has been used in economics for price analysis. The added complications of computation have caused this to be less popular than either of the other averages. although it is also capable of algebraic manipulation.

The mode is not affected by the size of the variations from the average, only the distribution. It is the location of greatest concentration and indicates a typical value for a reasonably large sample. With an unordered set of data, the mode is time consuming to locate and is not capable of algebraic manipulation.

16

The median is most useful when the center of an incomplete set is needed. It is not affected by extreme variations and is simple to find if the number of data points are known. Although it has some arithmetic properties, it is not readily adaptable to computational methods.

DISPERSION AND SKEWNESS

The center or central tendency of a data series is not a sufficient description for price analysis. The manner in which it is scattered about a given point, its dispersion and shewness, are necessary to describe the data. The mean deviation is a basic method for measuring distribution and may be calculated about any measure of central location, for example, the arithmetic mean. It is found by computing

$$MD = \frac{\sum \left| \text{price}_i - \frac{\sum \text{price}_i}{n} \right|}{n} \quad \text{or} \quad @Avg(@AbsValue(\text{price} - @Avg(\text{price}, n)), n)$$

where MD is the mean deviation, the average of the differences between each price and the arithmetic mean of the prices, or other measure of central location, with signs ignored.

The standard deviation is a special form of measuring average deviation from the mean, which uses the rootmeansquare

$$\begin{aligned} \text{Standard deviation} &= \sqrt{\frac{\sum \left(\text{price}_i - \frac{\sum \text{price}_i}{n} \right)^2}{n}} \quad \text{or} \\ &= @Sqrt(@Avg((\text{price} - @Avg(\text{price}, n))^2, n)) \end{aligned}$$

where the differences between the individual prices and the mean are squared to emphasize the significance of extreme values, and then total final value is scaled back using the square root function. This popular measure, found throughout this book, is available in all spreadsheets and software programs as @Std or @Stdev For n prices, the standard deviation is simply @Std(price,n).

The standard deviation is the most popular way of measuring the degree of dispersion of the data. The value of one standard deviation about the mean represents a clustering of about 68% of the data, two standard deviations from the mean include 95.5% of all data, and three standard deviations encompass 99.7%, nearly all the data. These values represent the groupings of a perfectly normal set of data, shown in Figure 24.

Probability of Achieving a Return


If we look at Figure 24 as the annual returns for the stock market over the past 50 years, then the mean is about 8% and one standard deviation is 16%. In any one year we can expect the compounded rate of return to be 8%; however, there is a 32% chance that it will be either greater than 24% (mean plus one standard deviation) or less than 8% (the mean minus one standard deviation). If you would like to know the probability of a return of 20% or greater, you must first rescale the values,

$$\text{Probability of reaching objective} = \frac{\text{objective} - \text{mean}}{\text{standard deviation}}$$

If your objective is 20%, we calculate

$$\text{Probability} = \frac{20\% - 8\%}{16\%} = .75$$

FIGURE 24 Normal distribution showing the percentage area included within one standard deviation about the arithmetic mean.



We look in Appendix A1 under the probability for normal curves, and find that a standard deviation of .75 gives 27.34%, a grouping of 54.68% of the data. That leaves onehalf of the remaining data, or 22.66%, above the target of 20%.

Skewness

Most price data, however, are not normally distributed. For physical commodities, such as gold, grains, and interest rates (yield), prices tend to spend more time at low levels and much less time at extreme highs; while gold peaked at \$800 per ounce for one day, it has remained between \$375 and \$400 per ounce for most of the past 10 years. The possibility of failing below \$400 by the same amount as its rise to \$800 is impossible, unless you believe that gold can go to zero. This relationship of price versus time, in which markets spend more time at lower levels, can be measured as skewness—the amount of distortion from a symmetric shape that makes the curve appear to be short on one side and extended on the other. In a perfectly normal distribution, the median and mode coincide. As prices become extremely high, which often happens for short intervals of time, the mean will show the greatest change and the mode will show the least. The difference between the mean and the mode, adjusted for dispersion using the standard deviation of the distribution, gives a good measure of skewness

$$(\text{Skewness}) S_K = \frac{\text{mean} - \text{mode}}{\text{standard deviation}}$$

Because the distance between the mean and the mode, in a moderately skewed distribution, is three times the distance between the mean and the median, the relationship can also be written as:

$$(\text{Skewness}) S_K = \frac{3 \times (\text{mean} - \text{median})}{\text{standard deviation}}$$

This last formula may be more practical for computer applications, because the mode requires dividing the data into groups and counting the number of occurrences in each bar. When interpreting the value of S_K , the distribution leans to the right when S_K is positive (the mean is greater than the median), and it is skewed left when S_K is negative.

Kurtosis

One last measurement, that of kurtosis, should be familiar to analysts. Kurtosis is the "peakedness" of a distribution, the analysis of "central tendency." For most cases a smaller standard deviation means that prices are clustered closer together; however, this does not always describe

the distribution clearly. Because so much of identifying a trend comes down to deciding whether a price change is normal or likely to be a leading indicator of a new direction, deciding whether prices are closely grouped or broadly distributed may be useful. Kurtosis measures the height of the distribution.

Transformations

The skewness of a data series can sometimes be corrected using a transformation on the data. Price data may be skewed in a specific pattern. For example, if there are 1/4 of the occurrences at twice the price and 1/9 of the occurrences at three times the price, the original data can be transformed into a normal distribution by taking the square root of each data item. The characteristics of price data often show a logarithmic, power, or square-root relationship.


Skewness in Price Distributions

Because the lower price levels of most commodities are determined by production costs, price distributions show a clear boundary of resistance in that direction. At the high levels, prices can have a very long tail of low frequency. Figure 25 shows the change in the distribution of prices as the mean price (over shorter intervals) changes. This pattern indicates that a normal distribution is not appropriate for commodity prices, and that a log distribution would only apply to overall long-term distributions.

Choosing between Frequency Distribution and Standard Deviation

You should note that it is more likely that unreliable probabilities will result from using too little data than from the choice of method. For example, we might choose to look at the distribution of one month of daily data, about 23 days; however, it is not much of a sample. The price or equity changes being measured might be completely different during the next month. Even the most recent five years of S&P data will not show a drop as large as October 1987.

FIGURE 25 Changing distribution at different price levels. A, B, and C are increasing mean values of three shorterterm distributions.



Although we can identify and measure skewness, it is difficult to get meaningful probabilities using a standard deviation taken on very distorted distributions. It is simpler to use a frequency distribution for data with long tails on one side and truncated results on the other. To find the likelihood of returns using a trend system with a stoploss, you can simply sort the data in ascending order using a spreadsheet, then count from each

end to find the extremes. You will notice that the largest 10% of the profits cover a wide range, while the largest 10% of the losses is clustered together.

A standard deviation is very helpful for giving some indication that a price move, larger than any we have seen in the data, is possible. Because it assumes a normally shaped curve, a large clustering of data toward one end will force the curve to extend further. Although the usefulness of the exact probabilities is questionable, there is no doubt that, given enough time, we will see price moves, profits, and losses that are larger than we have seen in the past.

Student test

Throughout the development and testing of a trading system, we want to know if the results we are seeing are as expected. The answer will keep referring back to the size of the sample and the amount of variance that is typical of the data during this period. Readers are encouraged to refer to other sections in the book on sample error and chisquare test. Another popular method for measuring whether the average price of the data is significantly different from zero, that is, if there is an underlying trend bias or if the pattern exhibits random qualities, is the student test,

$$t = \frac{\text{average of price changes}}{\text{standard deviation of price changes}} \times \sqrt{\text{No. of data items}}$$

and where degrees of freedom = number of data items - 1. The more trades in the sample, the more reliable the results. The values of t needed to be significant can be found in Appendix 1, Table A1.2, "T-Distribution." The column headed ".10" gives the 90% confidence level, ".05" is 95%, and ".005" is 99.5% confidence.

If we separate data into two periods and compare the average of the two periods for consistency, we can decide whether the data has changed significantly. This is done with a 2sample test:

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{\frac{v_1^2}{n_1} + \frac{v_2^2}{n_2}}}$$

where $\bar{x}_1$ and $\bar{x}_2$ are the averages of data periods 1 and 2, and n_1 and n_2 are the number of data items in periods 1 and 2, and df , needed to find the confidence level.

$$df = \frac{\left(\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2} \right)^2}{\frac{\left(\frac{s_1^2}{n_1} \right)^2}{(n_1 - 1)} + \frac{\left(\frac{s_2^2}{n_2} \right)^2}{(n_2 - 1)}}$$

The student test can also be used to compare the profits and losses generated by a trading system to show that the underlying system process is sound. Simply replace the data items by the average profit or loss of the system, the number of data items by the number

of trades, and calculate all other values using the profit/loss to get the student test value for the trading performance.

STANDARDIZING RETURNS AND RISK

To compare one trading method with another, it is necessary to standardize both the tests and the measurements used for evaluation. If one system has total returns of 50% and the other of 250%, we cannot decide which is best unless we know the duration of the test. If the 50% return was over 1 year and the 250% return over 10 years, then the first one is best. Similarly, the return relative to the risk is crucial to performance as will be discussed in Chapter 21 ("Testing"). For now it is only important that returns and risk be annualized or standardized to make comparisons valid.

Calculating Returns

The calculation of rate of return is essential for assessing performance as well as for many arbitrage situations. In its simplest form, the oneperiod rate of return R , or the holding period rate of return is

$$\text{Return} = \frac{\text{Ending value} - \text{Starting value}}{\text{Starting value}} = \frac{\text{Ending value}}{\text{Starting value}} - 1$$

$$R = \frac{P_1 - P_0}{P_0} = \frac{P_1}{P_0} - 1$$

where P_0 is the initial investment or starting value, and P_1 is the value of the investment after one period. In most cases, it is desirable to standardize the returns by annualizing. This is particularly helpful when comparing two sets of test results, in which each covers a different time period. Although calculations on government instruments use a 360day rate (based on 90day quarters), a 365day rate is common for most other purposes. The following formulas show 365 days; however, 360 may be substituted.

The annualized rate of return on a simpleinterest basis for an investment over n days is

$$R_{\text{simple}} = \frac{P_1 - P_0}{P_0} \times \frac{365}{n}$$

The *annualized compounded rate of return* is

$$R_{\text{compounded}} = \left(1 + \frac{P_n - P_0}{P_0} \right)^{(365/n)}$$

The geometric mean is the basis for the compounded growth associated with interest rates. If the initial investment is \$1,000 (P_0) and the ending value is \$1,600 (P_y) after 12 years ($y = 12$), there has been an increase of 60%. The simple rate of return is 5%, but the compounded growth shows

$$\text{Ending value} = \text{Starting value} \times (1 + \text{Compounded return})^y$$

$$P_y = P_0 \times (1 + R)^y$$

or

$$R = \left(\frac{P_y}{P_0} \right)^{(1/y)} - 1$$

$$R = \left(\frac{1600}{1000} \right)^{(1/12)} - 1$$

$$R = .0398 \text{ or about } 4.0\%$$

STANDARDIZING RETURNS AND RISK

The use of the standard deviation and compound probability of a return objective. In the following continuous returns is $\ln(1 + R_g)$, and it is assumed that

$$z = \frac{\ln\left(\frac{T}{B}\right) - \ln(1 + R_g)n}{\sqrt{s \times n}}$$

where z = standardized variable (can be looked up in Appendix A1)

T = target value or rate-of-return objective

B = beginning investment value

R_g = geometric average of periodic returns

n = number of periods

s = standard deviation of the logarithms of the quantities 1 plus the periodic returns

Indexing Returns

The Federal Government has defined standards for calculating returns in the Futures Industry Commodity Trading Advisors (CTAs). This is simply an indexing of returns based on the current period percentage change in equity. It is the same process as creating any index, and it allows trading returns to be compared with, for example, the S&P Index or the Lehman Brothers Treasury Index, on equal footing. Readers should refer to the section later in this chapter, "Constructing an Index."

Calculating Risk

Although we would always like to think about returns, it is even more important to be able to assess risk. With that in mind, there are two types of risk that are important for very different reasons. The first is catastrophic risk, which will cause fatal losses or ruin. This is a complicated type of risk, because it may be the result of a single price shock or a steady deterioration of equity by being overleveraged. This form of risk will be discussed in detail later in the book.

Standard risk measurements are useful for comparing the performance of two systems and for understanding how someone else might evaluate your own equity profile. The simplest estimate of risk is the variance of equity over a time interval commonly used by most investment managers. To calculate the variance, it is first necessary to find the mean return, or the expected return, on an investment:

$$E(R) = p_1R_1 + p_2R_2 + \cdots + p_nR_n$$

The most common measure of risk is variance, calculated by squaring the deviation of each return from the mean, then multiplying each value by its

associated probability

$$\mathbf{variance} = \sum_{i=1}^n p_i R_i$$

The sum of these values is called the variance, and the square root of the variance is called the standard deviation. This was given in another form in the early section "Dispersion and Skewness."

$$\mathbf{standard\ deviation} = \sqrt{p_1[R_1 - E(R)]^2 + p_2[R_2 - E(R)]^2 + \cdots + p_n[R_n - E(R)]^2}$$

$$= \sqrt{\sum_{i=1}^n p_i [R_i - E(R)]^2}$$

This and other very clear explanations of returns can be found in Peter L. Bernstein's *The Portable MBA in Investment* (John Wiley & Sons, New York, 1995).